15. R&D Plan

15.1 Introduction

In this section we summarize the key R&D activities required to validate the design concepts described in this Neutrino Source Feasibility Study. The items covered here fall into two categories: *i*) those required to validate or improve the components that drive the fabrication costs of the facility, and *ii*) those required to address the performance and/or feasibility of fabrication of particular components. In the first case, R&D will mainly involve hardware fabrication and testing without beam. In the second case, performance tests with beam will generally be required in addition to prototyping. Much of the hardware R&D of this second class must be guided by knowledge gleaned from simulation studies in which realistic errors on component performance are included.

What is presented here a long-term schedule for carrying out the R&D activities. A technology-limited schedule in which we assume that there is available staff to carry out the program is assumed, and available funding to support it. With these assumptions, we believe that, at the end of this period, the R&D program will have progressed to the stage where we could confidently begin developing a Conceptual Design Report for a Neutrino Factory. It is important to note that much of the hardware development effort envisioned should proceed in parallel on several fronts and at different laboratories. Indeed, it *must* proceed this way if we are to complete the R&D tasks in a reasonable time frame. Clearly, however, our progress on this time schedule requires funding commensurate with the program needs; this is the resource over which we have the least control.

15.2 Proton Driver

The main technical issue to address here is the production of high intensity, short proton bunches, with a beam power of the order of 1 MW or more and with very small beam loss during acceleration. We are fortunate that there are a number of existing proton synchrotrons and linacs available for testing some of the concepts and components needed. This program is already under way at several laboratories, not all of which are motivated by the need for an intense muon beam.

15.2.1 Hardware Development

Hardware that may be required for this purpose in the synchrotron includes high-gradient, low frequency RF cavities to produce short bunches, inductive inserts to compensate for space-charge blowup, a thin-walled chamber to minimize eddy current effects, and tracking circuitry to maintain the ratio between various magnet families during the ramp cycle. Development of effective collimator systems, feedback systems, fast full-aperture kickers, and efficient and reliable power supply technology are also important R&D topics. Hardware for the initial linac includes an RF chopper, a high-brightness proton source, and a high-current radio-frequency quadrupole (RFQ).

15.2.2 Experimental Program

Experiments to study intense short, bunches have started and will continue. These will involve using inductive inserts, testing the behavior of RF cavities with beam (beam loading compensation, impedance effects), and studies of the microwave instability below transition energy. Tests of a cavity loaded with Finemet will be carried out to study its behavior in a beam. These will initially be done with a CW system at a gradient of about 30 kV/m and a frequency of a few MHz. Studies of much higher cavities are underway. Bunch compression tests to study methods of reaching high peak current and nanosecond bunch lengths will be done at a number of laboratories (BNL, CERN, Fermilab, GSI, Indiana University, and KEK).

15.3 Target System

The primary technical issue to address in this area is the survivability of a target subjected to a proton beam power of about 1,5 MW. Both solid and liquid (mercury-jet) targets are candidates for the Neutrino Source. These pose different problems, but experimental validation of the available computer codes is needed in either case. In the context of this Study, R&D for target systems is a broad category that includes development work for the mechanical design of the target and beam dump, design of the remote handling hardware, and safety and environment issues that may impact the target system design. These are divided into near-term activities that should begin next fiscal year and long-term activities for the next two to five years. General issues related to target performance are mainly covered under the Targetry R&D program of the Muon Collaboration. Included in this broader category are: the calculational and experimental studies of the behavior of a mercury-jet target; yield measurements from the target (both pions, which

determine the muon yield, and neutrons, which define the radiation environment near the target); lifetime and power deposition tests with a carbon target; materials studies for the superconducting 20-T target solenoid; studies of the proton beam dump; and exploration of solid-target geometry and cooling options, such as a band target configuration.

15.3.1   Near-Term R&D

The primary short-term goal is to determine the survivability of a passively cooled graphite target under the thermal and radiation conditions expected at a Neutrino Factory. A test of a strained graphite target is planned at the 1-MW Spallation Radiation Damage Facility at LANL very soon in collaboration with colleagues from ORNL. This will involve the following activities: *i*) evaluation of the properties of different grades of graphite, including carbon-carbon composites; *ii*) thermal tests to determine bending-creep and axial-creep effects on full-scale target rods for various support schemes; *iii*) investigation of non-uniform power distribution on full-scale rods to determine temperature and stress effects; *iv*) thermal shock tests on small scale samples; *v*) neutron irradiation tests to determine target survivability; *vi*) designing and testing of one or more water cooling schemes for the target support tube.   Other activities will include investigation (by calculation and tests) of the pressure shock-wave effects due to the very short duration beam spill, development of a design for a water-cooled graphite target if tests on the passively cooled target are unsuccessful (i.e., fabrication of small-scale samples of graphite with cooling tubes and testing them for thermal and radiation survivability), and evaluation of safety and environmental requirements. Beam tests at the A3 line of the AGS at BNL will be carried out, beginning next year.

15.3.2   Long-Term R&D

After the initial tests, the next step will be to design and fabricate a full-scale prototype of a target and its support tube, complete with coolant interfaces. This target will be tested in a proton beam under realistic conditions. Another important task will be to develop design details for the high-power beam dump, which is integrated in the decay channel. It is envisioned that we will carry out small-scale tests to demonstrate feasibility. Work on the alternative of a mercury-jet target will also be pursued. This will include operation of a jet target in a high magnetic field at the NHMFL. Beam tests at the AGS will continue.   Ultimately, yields will be measured.

 The most challenging magnet is the capture solenoid in the target region, and clearly the most critical area in that subsystem is the resistive poly-Bitter magnet nearest the target. This magnet is critical primarily because the 20 T magnetic field cannot be achieved without it. Magnet lifetime is expected to be dominated by erosion associated with the extreme cooling requirements but, in addition, the device has an uncertain radiation tolerance. There is a large body of data on the radiation tolerance of a wide range of insulating materials and there is also very good experience with the poly-Bitter technology at the NHMFL. While there is every reason to expect that a good solution can be found, it is important that a careful study be undertaken to unify existing experience, to incorporate candidate materials into a systematic design process, and to demonstrate compatibility of candidate materials and components in a working poly-Bitter magnet model. This may require beam time at the 1-MW Spallation Radiation Damage Facility at LANL.

There are also "facility" design issues worthy of R&D effort. We will develop full-scale mock-ups (using simulated components and substitute materials where possible) to demonstrate remote handling of key components. A partial mock up of the target region, complete with utility interfaces (electrical, cooling, instrumentation) will be fabricated. This setup will be used to demonstrate replacement of the target module, replacement of the Bitter coil, and replacement of the proton beam window. Finally, consequences of possible failure modes, e.g., of the beam isolation window, will be investigated to make sure the design is robust against such occurrences. We expect to benefit in this area from close collaboration with groups at existing facilities (or those now under construction, for example the SNS at ORNL) that have similar requirements for beam power on target.

15.4   Decay, Capture, and Bunching

One of the more challenging areas of the front-end system is the phase rotation section, which makes use of an induction linac system to reduce the energy spread of the beam prior to bunching and cooling it. This device is intended to operate at an accelerating gradient of 2 MV/m (from –0.5 to +1.5 MV/m) with a pulse structure consisting of four ~150 ns pulses within a 2 μs interval.   In addition, it must accommodate superconducting solenoid coils in the accelerating gaps to maintain the transverse focusing of the beam along the length of the linac. The design considered here is based on a device of similar size, but lower gradient and normal conducting coils for focusing, built for the DARHT project. Nonetheless, the combination of high gradient, internal magnetic field, and untested pulse structure mean that a prototype induction cell and pulser must be built and tested. The R&D program will determine whether the induction linac can provide the acceleration waveforms needed for the phase rotation. Important issues are the accuracy

and reproducibility of the voltage waveform, determining limits on the gradient, and the lifetime and reliability of the modules. The integration of the internal superconducting solenoids with the accelerating structure must be tested.

As discussed in Chapter 5, the beam energy spread after phase-energy rotation and bunching is larger than desired, even in the induction linac phase rotation scenario. A smaller energy spread may be accomplished by better optimization of the decay section and induction linac, or it may make use of an initial phase rotation using RF cavities. This is an area that will benefit from more integrated simulation studies of the complete cooling system.

15.4.1   Hardware Development

The induction linac R&D effort will include the design, construction, and testing of a prototype module of the 200 MeV induction linac. This module will consist of a single induction cell (–0.5 to 1.5 MV), a vacuum chamber, a four-pulse generator module, and a superconducting magnet assembly with its associated cryogenic cooling system and power supply. Because the induction cores provide the dominant load for the four-pulse generator module, no beam test is necessary.

Initial activities will consist of conceptual design work and small-scale component tests. Design tradeoffs will be performed to establish a set of beam parameters, such as pulse width vs. pulse separation along the linac, that are consistent with pulse-generator constraints. Further study will be done to optimize the design of the pulse-train generator. Specifically, recent and anticipated advances in the technology of high voltage/high power switches will be explored and samples tested in an attempt to reduce or eliminate the magnetic compressor pulse charge stages. Alternate pulse generation topologies will be explored and a detailed circuit simulation model will be developed and exercised. Induction cell core tests will be performed to establish accurate loss current vs. time characteristics. A full-scale core using the most promising material will be built and tested. Engineering will be done on the superconducting magnet to bring it to the conceptual design level.

After defining the system, we will fully prototype such a cell and its pulser. First, hardware for a 1 m cell and a pulser capable of two-pulse operation will be procured. Initial testing will begin, first in a single-pulse mode and then in a two-pulse mode. Hardware will next be procured for an additional two pulses, after which testing of the full four-pulse prototype system will commence. This staged approach to testing and hardware procurement allows early identification of problems and permits their solutions to be incorporated into the design before committing to all the hardware. Parallel simulation efforts will finalize the voltage profile requirements that the pulser and induction linac must meet.

A number of different solenoid magnet designs are needed for the capture and decay channel region. Though most of these are straightforward, it will be prudent to develop prototypes of each type to ensure that there are no lurking fabrication issues. Throughout this region, it will be necessary to develop and optimize both cooling and quench protection schemes, with an eye toward enhancing system reliability.

The buncher requires cavities with lower gradient than those of the cooling channel, and it may be possible to enlarge their aperture. If simulation studies confirm this, development of a modified cavity will be needed. If a two-frequency buncher is adopted, a second-harmonic cavity will be needed. Ideally, this cavity would maintain the same aperture as the lower frequency cells, which makes its aperture relatively large.

15.5   Cooling Channel

15.5.1   Cooling Theory, Simulation, and Optimization

Although considerable progress has been made in understanding the issues involved in applying ionization cooling and in designing realistic cooling channels, more simulation work is clearly needed to optimize performance, to test error sensitivity, and to minimize cost. In particular, the present baseline design does not fully meet the requirements for intensity at the end of the cooling channel. Cooling the longitudinal phase space would ease the design requirements for all downstream components (acceleration and storage ring portions of the facility), thus reducing cost and permitting a more efficient implementation. The current strategy is to examine emittance exchange (longitudinal-to-transverse) because, based on our present simulations, we are confident that this would increase transmission in the cooling channel. An in-depth look at this question will occur in Fall '00 when a planned Emittance Exchange Workshop takes place at BNL. A solenoid channel with only a single polarity change appears promising in initial simulations, and we will follow up on this concept. In addition, we wish to consider a configuration based on a helical magnetic channel. Once a scheme is developed, we need to understand the uncertainties in its expected performance with respect to various model assumptions (e.g., correlations between straggling and large-angle-scattering, the atomic form factor in models used in

estimating the multiple scattering distribution and energy straggling, and the effect of intense magnetic fields on these processes). Work will continue to refine the figure of merit that characterizes the front-end performance. In addition, particle distributions at the end of the cooling channel will be transmitted through the acceleration and storage ring sections to ensure a self-consistent design for the entire complex.

In Section 6 we presented two ionization cooling schemes based on differing design principles. The objective of our study is to design an efficient transverse cooling channel with components that can be built at the present time, or could be developed with a well-defined R&D plan. Continued interaction and iteration with the engineers will be a key part of our work. Detailed simulation studies are required to obtain the optimal solution for each of the available cooling channels. These studies will be crucial in selecting the best design, based on performance, engineering constraints, and cost. For example, the engineering details of the RF-cavity grids or windows (see Section 6), and their corresponding electric fields, will be implemented in our simulations. This activity will require the use of full three-dimensional codes and diagnostics to estimate reliably both the effects of scattering from these complicated shapes and the effects of the field distortions on the longitudinal phase space. Evaluating the effect of alignment errors also requires three-dimensional codes. Another activity that requires code development is the study of cooling channels that can achieve longitudinal cooling, because these involve elliptical or wedge-shaped absorbers. In addition, the optimization of each design must respect the engineering constraints on coil current densities, coil placement, and forces on the coils. Successive iterations of our simulation studies with the engineering analysis of each of the design variants will achieve the best solution.

The performance of any cooling channel is tightly coupled to the performance of the machine components upstream of it and to the acceptance of the acceleration section that follows. For this reason, the optimization of the cooling channel must be iterated in the context of the designs of these other sections. Based on the present work, there are several places where additional effort is needed. The transverse matching between the buncher section and the downstream cooling channel needs to be improved, and the longitudinal bunching itself can be made somewhat more efficient. The combination of the induction linac and its upstream drift section needs to be optimized to reduce the energy spread of the beam going into the cooling channel. The cooling channel itself needs further optimization. The figure of merit used is the number of muons per incident proton within a specified transverse and longitudinal emittance. For now, we choose the transverse emittance cut to be 6000 mm-mrad, i.e., four times the rms emittance initially specified. This value can easily be accepted in the acceleration section to follow.

To provide a realistic experimental validation of the cooling concept we adopt, an experimental program (MUCOOL) has been initiated at Fermilab. This effort will design, fabricate, and test all of the components required for a cooling channel cell—absorber, RF cavity (both a beryllium foil and a gridded cavity will be built and tested), and superconducting solenoid. This program will focus initially on producing prototypes of all these components and bench testing them. (As noted in Section 15.5.2, the absorber will be tested in a beam of protons or electrons as part of its initial commissioning.) After all components are available, a complete cooling cell will be assembled and its performance tested cryogenically and electrically. Part of the MUCOOL task is to develop instrumentation capable of measuring the muon beam properties within the cooling channel. Access to a beam will be necessary for this work.

15.5.2   Absorbers

The liquid hydrogen (LH_2) cells that will serve as the cooling channel absorbers have a number of R&D topics that need study. The baseline (FOFO) cooling design requires liquid-hydrogen absorbers that are thin (~13 cm) relative to their diameter (20–30 cm). For a given pressure differential, hemispherical windows are thinnest but, with this oblate shape, thicker ellipsoidal or torispherical windows are required to provide a sufficiently short sagitta. A program of design studies backed up by carefully designed tests will be needed to establish safe design and operating parameters for the LH_2 containment windows. Because multiple scattering in the windows is a major source of "heating" of the muon beam, every effort must be made to minimize window thickness. Aluminum is our default material choice, but AlBeMet—a beryllium-aluminum alloy—is an attractive alternative if it is shown to be compatible with liquid hydrogen. With 40% greater strength than aluminum and 2.1 times the radiation length, AlBeMet has the potential to lower the total radiation-length fraction per absorber from ~2.4% to 1.8% or less, depending on the detailed optimization of absorber dimensions. (While beryllium windows may also be feasible, there appears to be little additional gain compared with AlBeMet.)

In all scenarios the specific power dissipation in the absorbers is large and represents a significant cryogenic load. Handling this heat load is a significant design challenge in terms of fluid dynamics and requires sophisticated thermal modeling of target heating. An R&D program is now under way at the Illinois Institute of Technology (IIT) to understand the thermal and fluid-flow aspects of maintaining a constant temperature within the absorber volume despite

the large spatial and temporal variations in power density. This program is beginning with computational-fluid-dynamics studies and is planned to proceed to bench tests and high-power beam tests of an absorber prototype over the next year. The prototype absorber will be built and initially tested for safety and then with a proton or electron beam to explore pulsed heating effects. Ultimately, one or more absorbers of a selected design will be tested.

In some scenarios (especially those with emittance exchange), lithium hydride (LiH) absorbers may be called for. Since it is a solid, LiH can in principle be fabricated in arbitrary shapes. In emittance-exchange channels, dispersion in the lattice spatially separates muons according to their energies, whereupon specially shaped absorbers can be used to absorb more energy from muons of higher energy and less from those of low energy. Unfortunately, solid LiH shapes are not commercially available, and procedures for their fabrication need to be developed. Such an effort is challenging since LiH is very reactive.

15.5.3   Diagnostics

Techniques for optimizing the operation of a physical cooling channel must also be developed. The beam emittances and particle losses in these cooling channels must be measured in order to optimize running conditions. These beam measurements will be complicated by the large size of the beam, the limited space available for detectors, the high magnetic fields, and the need for low-temperature insulation. Although measurements of muon beams have been done for years, the issues associated with the cooling channel will make the required measurements difficult. For this reason, some R&D directed at beam diagnostics and beam measurements is clearly required for the design of the neutrino source.

Diagnostic information for the following is needed: *i*) initial matching of the cooling optics to the beam parameters, *ii*) accuracy of this match down the length of the cooling channel, *iii*) the accuracy of physical alignment of beam components, *iv*) identification of transverse and longitudinal loss mechanisms, and *v*) measurement of the emittance at various stages of cooling. Because the emittance will only change by only a few percent in each cooling cell, it will be desirable to have a few special diagnostic sections interspersed with cooling sections to make precise measurements.

This is an R&D area in which University groups can make significant contributions because new ideas are required in this area. Developing concepts as well prototypes perfectly fits infrastructures that are available in Universities while actual beam tests would be done in close collaboration with the participating laboratories.

15.6   Acceleration

There are a number of technical issues in this area that will require R&D effort. A very short section of linac that immediately follows the cooling channel might use normal conducting RF, i.e., it is basically a continuation of the cooling channel RF system, without the absorbers. However, it will be followed by a superconducting linac operating at the same frequency (201 MHz) to raise the beam energy to 3 GeV. The 201 MHz SCRF linac cavities must be designed and prototyped, and must be compatible with the attendant magnetic focusing system. Even weak stray magnetic fields from the focusing system would render the superconducting RF (SCRF) cavities inoperative, so effective magnetic shielding is a necessity. These cavities give rise to several challenges. Although we envision filling them with a low-power, long-pulse technique that reduces the number of klystrons and diminishes the power requirements on the input coupler, the Lorentz force detuning and microphonics remain significant issues. These effects are at least an order of magnitude more severe than in comparable installations at CERN, JLab, SNS, or CESR. Given the critical role played by the 201-MHz SCRF in our design, these issues merit special attention. Close collaboration with universities and institutes in the US as well as abroad is foreseen. We plan to encourage and support the technical development of the cavity structures as well as the testing. We have identified the acceleration as one of the major cost drivers and therefore have to have a strong program in this area very soon to reduce costs.

Optimization of the RLAs in terms of the number of passes in each must be done. This depends on the details of the splitter and recombiner magnet designs and also on the beam energy spread coming from the cooling channel. Designs for these magnets must be developed and—depending on how nonstandard they are—prototypes will be needed. Work on finalizing the optics design for the arcs must be done. Designs that avoid the use of strong sextupoles look promising but must be simulated based on the actual particle distribution provided from the front-end simulation effort. Optimization of the arcs in terms of magnetic fields and technology (room temperature or superconducting magnets) must be completed. With SC magnets, radiation heating becomes an issue and must be assessed and dealt with. Assessment of field error effects on the beam transport must be determined to define acceptance criteria for the magnets. This will require use of sophisticated tracking codes like COSY that permit rigorous treatment of field errors and fringe-field effects. Because the beam circulates in each RLA for only a few turns, the sensitivity to magnet errors should not

be extreme, though the large energy spread will enhance the effects. Though it seems likely that the RF cavity frequency in RLA1 will be 201 MHz, the RLA2 frequency could be either 201 or 402 MHz. The choice of the higher frequency has been made provisionally but should be validated by performance simulations. In particular, we must develop a longitudinal matching strategy that manages the single-bunch phase space through the acceleration cycle and we must verify that the increased energy droop associated with the higher frequency cavities can be accommodated in the beam transport. While a preliminary estimate of the beam-breakup instability shows that it should not be a problem, a full assessment of collective effects is needed in each RLA to determine whether higher-order-mode dampers are needed for the RF cavities.

15.6.1   Hardware Development

A prototype RF cavity will be built and tested at gradients up to 15 MV/m with a Q of $\approx 5 \times 10^{9}$. This cavity must also operate in a channel having decaying muons. Though SCRF cavities at higher frequencies are by now becoming common, cavities sized for 201 MHz have never been built. The large size of the cavity makes it more sensitive to microphonics; a prototype cavity will permit us to assess this aspect. Even extending fabrication and cleaning techniques to this size range is an R&D topic in itself. Developing the proper facilities is a long lead-time item and will require significant up-front expenditures. At high gradients, the input power coupler design will need to be tested and validated, though the chosen slow-fill technique is expected to make this straightforward. Detuning issues associated with the pulsed RF system must be evaluated. Finally, because of the large stored energy in a 201 MHz SCRF cavity, a reliable quench protection system must be designed and tested.

Prototypes of the special magnets (splitters, recombiners, and kickers) will be designed. If curved SC dipoles are used, a prototype will be fully engineered. Effective magnetic shielding concepts must be developed to protect the SCRF cavities. Beam diagnostics that function properly in a beam pipe containing a decaying muon beam must be designed and prototypes built and tested.

15.7   Storage Ring

The most serious R&D issue for the storage ring operation is a more exact understanding and modeling of large-aperture, strong quadrupole fringe fields. Once the end fields are accurately described (which appears to require detailed magnet designs), they must be inserted into the tracking programs and carefully studied by themselves and in combination with field and alignment errors.

Also, the requirements on beam steering in the production straight section, intensity monitoring, collimation of electrons from the long straight section etc have to be specified to see if they can be achieved.

15.8   Solenoids for Decay, Phase Rotation, and Cooling Channels

Although design and implementation of every solenoid channel discussed above is a nonstandard technical task, one can expect to meet the most challenging problems working with the cooling channel. Both radial and axial ponderomotive forces in this channel are extremely large as well as the interacting forces acting on the first and the last magnets in the string. Because the mechanical behavior of a magnet in this channel is the main issue to study, thorough modeling of the system is a must in this case. Other objects of study are optimization of the quench protection system and analysis of stresses in coils during cooling down and warming up. Undoubtedly, modeling and prototyping of the channel magnets must accompany development of magnetic systems for the Neutrino Factory front end.

The R&D scenario suggested for the next step includes:

- Optimization of the magnet design for the cooling channel with the goal of reducing mechanical stress in coils;
- Mechanical stability analysis for the magnets of the cooling channel, development of the channel mechanical scheme, and design of a mechanical support for coils and magnets in the cooling channel;
- Development of Nb_3Sn cable to use in the cooling channel; development of coil fabrication technology;
- Development and optimization of the quench protection system and the cooling scheme for each channel with the goal of increasing system reliability and reducing the time for cooling down and warming up;
- Building full-scale prototypes of magnets in each channel.

15.9   Cryogenic Systems

The estimates presented above place the cryogenic system of the Neutrino Factory within current practice, and no R&D is required. The next step to take is to try to get concepts for all of the cryogenic modules of the Neutrino Factory and

to assemble as much design detail as possible. It is important also to begin to think about operating scenarios. With this additional information a much better estimate of the scope of the cryogenic system can be made, some assessment of its operating requirements can be developed, and we can begin to provide support for the component design process.

15.10   High Power RF Sources and Normal Conducting Cavities

Operation of the cooling channel requires normal conducting (NC) RF cavities operating at a frequency in the range of 175–201 MHz and a high gradient of about 15 MV/m. Moreover, the cavities are immersed in a strong solenoidal magnetic field. To increase the cavity shunt impedance (lowering power requirements) two alternatives are being explored. In one option, each cavity aperture is covered with a thin beryllium foil (125 μm thick) that serves to enhance the on-axis accelerating gradient while giving rise to only a modest amount of multiple scattering. In the second approach, the cavity apertures have a grid of tubes to achieve the same benefits. Both techniques have significant fabrication issues, and it is planned to build prototypes of both cavity types to test them. Either of these approaches causes some increase in multiple scattering, and thus some degradation of the cooling effect of the channel. Simulations will be used to evaluate the acceptability of each approach from the scattering viewpoint. High-power cavities built with each of the two candidate techniques will be tested for voltage handling capability, first without, then with a superimposed solenoidal magnetic field. This will permit evaluation of breakdown and multipactor behavior. In the beryllium foil case, heating from RF fields in the cavity (which scales as r^4) defines an upper limit for the aperture radius. Low-power tests of foil deflection are under way to validate the dimensions chosen; these must be repeated for the high-power test cavity. Fabrication techniques must be developed for attaching the beryllium foils to the cavity in a mechanically reliable manner while maintaining good thermal and electrical contact. High-power tests of the gridded cavity will be carried out to verify thermal behavior and electrical breakdown at high gradients. Fabrication of the grid tubes from AlBeMet rather than aluminum is attractive from a multiple scattering viewpoint and will be tested.

Development of a reliable high power ($\geq$ 10 MW) RF source is another important R&D activity for the cooling channel. Candidate technologies include multibeam klystrons, single-beam klystrons, eventually inductive output tubes (IOTs) or even lower-gain devices. We envision development of prototypes in close collaboration with the available expertise in other laboratories (e.g. SLAC) and industry. Multibeam klystrons seem to be the most promising approach at present because of their potential for efficiency and compactness. These devices at some point require industrial development, which in turn means long lead times and significant up-front costs. A compact solid-state modulator based in IGBT technology will also be build and tested. After coupling the RF power source to multicell superconducting RF cavities, we will assess the need to have individual frequency and phase control of each cell.

15.11   Summary

Future R&D on the different subsystems that have been described in this report will involve a strong and growing simulation effort as well as an accompanying hardware R&D program. In some places more patience is required to define the actual hardware R&D so that careful decisions can be used to drive the work in a direction that is useful. Bad decisions could delay progress on the hardware developments that are really required. For the cooling channel as a whole this is even more important than for its component parts. From this report and from R&D items described in this chapter the appropriate programs can be identified and launched as soon as possible with the funding presently available.

The R&D program which is described in this chapter does not have a schedule attached to it which would allow either an estimate of the time necessary to perform the work nor tracking the actual progress that is required. Although laboratories and universities which could perform the R&D work have been identified, the goals, the people to work on the R&D, the funding and many other critical issues have not been dealt with. These details will have to be negotiated and folded into a schedule before it is useful and can be presented.

This R&D chapter nevertheless circumscribes the cost drivers as well as the technically challenging problems. These cost drivers are not necessarily the performance limiting components in the Neutrino Factory and the performance limiting components do not always drive cost. Both nevertheless will have a strong influence on the time that is required to deal with these subjects thoroughly. Both will also drive the number of people that are required in order achieve these goals.

For a facility of the scale that is presented in this report the R&D program is very different from what we are used to in High Energy Physics. While the detailed development of single components that can be mass produced (one way or the other) for the final project is the usual goal, the R&D here spreads over a wider range of issues and technologies. It

requires the detailed development of many single components and for this reason needs much more support than usually anticipated and which can be spread over a large community. This support is what we are asking for.